

RIACS FINAL REPORT

January 1, 1996 through September 30, 1996

Contract NAS 2-13721

Submitted to:

Contracting Officer: Catharine P. Levin
Code IC, MS 241-1

FINAL

IN-61-CR

018915

Submitted by:

Research Institute for Advanced Computer Science (RIACS)

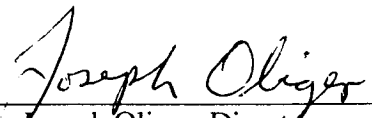
An Institute of:

Universities Space Research Association (USRA)

RIACS Principal Investigator
Joseph Oliger

NASA Technical Monitor
Henry Lum

COTR
Sonie Lau



Joseph Oliger, Director

TABLE OF CONTENTS

I. INTRODUCTION	3
II. RESEARCH PROJECTS	4
A. Advanced Methods for Scientific Computing	4
B. High Performance Networks	12
III. TECHNICAL REPORTS	17
IV. PUBLICATIONS	25
V. SEMINARS AND COLLOQUIA	27
VI. OTHER ACTIVITIES	27
VII. RIACS STAFF	29

I. INTRODUCTION

Joseph Oliger, Director

The Research Institute for Advanced Computer Science (RIACS) was established by the Universities Space Research Association (USRA) at the NASA Ames Research Center (ARC) on June 6, 1983. RIACS is privately operated by USRA, a consortium of universities with research programs in the aerospace sciences, under contract with NASA.

The primary mission of RIACS is to provide research and expertise in computer science and scientific computing to support the scientific missions of NASA ARC. The research carried out at RIACS must change its emphasis from year to year in response to NASA ARC's changing needs and technological opportunities. A flexible scientific staff is provided through a university faculty visitor program, a post doctoral program, and a student visitor program. Not only does this provide appropriate expertise but it also introduces scientists outside of NASA to NASA problems. A small group of core RIACS staff provides continuity and interacts with an ARC technical monitor and scientific advisory group to determine the RIACS mission. RIACS activities are reviewed and monitored by a USRA advisory council and ARC technical monitor.

Research at RIACS is currently being done in the following areas:

Advanced Methods for Scientific Computing
High Performance Networks

During this report period Professor Antony Jameson of Princeton University, Professor Wei-Pai Tang of the University of Waterloo, Professor Marsha Berger of New York University, Professor Tony Chan of UCLA, Associate Professor David Zingg of University of Toronto, Canada and Assistant Professor Andrew Sohn of New Jersey Institute of Technology have been visiting RIACS.

January 1, 1996 through September 30, 1996 RIACS had three staff scientists, four visiting scientists, one post-doctoral scientist, three consultants, two research associates and one research assistant.

RIACS held a joint workshop with Code I 29-30 July 1996. The workshop was held to discuss needs and opportunities in basic research in computer science in and for NASA applications. There were 14 talks given by NASA, industry and university scientists and three open discussion sessions. There were approximately fifty participants. A proceedings is being prepared. It is planned to have similar workshops on an annual basis.

RIACS technical reports are usually preprints of manuscripts that have been submitted to research journals or conference proceedings. A list of these reports for the period January 1, 1996 through September 30, 1996 is in the Reports and Abstracts section of this report.

II. RESEARCH PROJECTS

A. ADVANCED METHODS FOR SCIENTIFIC COMPUTING

DYNAMIC MESH ADAPTION OF TETRAHEDRAL GRIDS WITH QUALITY CONTROL

Rupak Biswas and Roger C. Strawn (US Army AFDD)

Dynamic mesh adaption on unstructured grids is a powerful tool for computing steady and unsteady three-dimensional problems that require local grid modifications to efficiently resolve solution features. By locally refining and coarsening the mesh to capture flowfield phenomena of interest, such a procedure makes standard computational methods more cost effective. However, for unsteady flows, the coarsening and refinement steps must be executed frequently; hence, their performance must be comparable to that of the flow solver.

An efficient solution-adaptive procedure has been developed for the simultaneous coarsening and refinement of three-dimensional tetrahedral meshes. An innovative data structure, that uses a combination of dynamically-allocated arrays and linked lists, allows the mesh connectivity to be rapidly reconstructed after individual points are added and/or deleted. The data structure is based on edges of the mesh rather than the tetrahedral elements that not only enhances the efficiency but also facilitates anisotropic mesh adaption. This means that each tetrahedral element is defined by its six edges rather than by its four vertices. Error indicators are used to identify regions of the mesh that require adaption. The overall objective is to optimize the distribution of mesh points so that the flowfield is accurately modeled with a minimum of computational resources.

One problem with the anisotropic subdivision of tetrahedral elements is that repeated refinement can significantly deteriorate the quality of the mesh. Poor mesh quality is defined as a grid deficiency that leads to inaccurate flowfield solutions. Our algorithm controls the quality of the mesh by never further subdividing anisotropically-refined elements. This effectively provides an upper bound on element face angles and controls the growth of the maximum vertex degree.

The tetrahedral mesh adaption code, called 3D-TAG, has been coupled with several unstructured flow solvers to solve large realistic problems on a Cray C-90. It has been used with a hybrid structured/unstructured overset grid scheme to capture vortices in a helicopter rotor wake. It has been combined with Euler solvers to solve problems in helicopter aerodynamics as well as to capture flows over complete supersonic transport configurations. The adaption procedure has also been successfully implemented on both a shared-memory SGI Power Challenge XL and a distributed-memory IBM SP2.

HELICOPTER NOISE PREDICTIONS USING COMPUTATIONAL AEROACOUSTICS

Rupak Biswas, Leonid Oliker and Roger C. Strawn (US Army AFDD)

High-performance helicopters and tiltrotors generate excessive noise both in forward flight and during takeoffs and landings. Future designs must have low noise if they are to operate successfully near heavily-populated areas. The accurate prediction of helicopter noise is essential to its control. Traditional acoustic analogy approaches cannot model the near-field nonlinear phenomena from high-speed rotor blades. CFD techniques are much better suited for these near-field nonlinearities but cannot efficiently propagate acoustic signals over large distances due to numerical dissipation. The Kirchhoff formulation, on the other hand, can propagate acoustic signals without dissipation but assumes a constant speed of sound outside a Kirchhoff surface. Our goal is to accurately and efficiently predict helicopter noise for a wide variety of flight regimes and rotor blade shapes using a hybrid CFD/Kirchhoff scheme.

Two different Kirchhoff methods have been developed to predict acoustic signals in the far field in a computationally efficient manner. Both methods perform an integration over a surface that completely encloses the rotor blades. The first method uses a surface that rotates with the blades. The second uses a nonrotating surface. Both the stationary and rotating Kirchhoff methods have been utilized to model high-speed impulsive (HSI) and blade-vortex interaction (BVI) noise for helicopter rotors in forward flight. The nonrotating method has also been combined with an overset grid Navier-Stokes solver to predict the acoustic signature of arbitrary rotors.

Up until now, these CFD/Kirchhoff techniques have been used to compute acoustic signals at a handful of far-field observer locations to compare with experimental microphone measurements. While these comparisons are useful for validation, they do not exploit the full capabilities of the new acoustic prediction methods. The CFD/Kirchhoff formulations can compute far-field acoustic pressures at many observer locations covering large regions of the flowfield. When viewed as a whole, these acoustic signals give a great deal of insight into the far-field propagation characteristics of helicopter noise.

HIGH-ORDER FINITE DIFFERENCE SCHEMES WITH SHARP SHOCK RESOLUTION FOR EULER EQUATIONS

Margot Gerritsen and Pelle Olsson

In previous research projects it has been demonstrated how to construct numerical schemes that support one- or two-point shocks for scalar conservation laws. This theory does not directly apply to the Euler equations, but following a similar approach we can construct a scalar viscosity that supports approximate one-point shocks. The viscosity is determined completely by the flow variables on either side of the shock.

The shock resolution can be significantly improved by the introduction of a small subgrid locally around the shock (steady as well as unsteady). The positioning of the subgrid is determined by a detection algorithm based on a multiscale wavelet analysis of the pressure grid function, which quickly and accurately locates potential shocks and spurious oscillations. It also supplies the information needed to compute the artificial viscosity terms. The detection algorithm is derived from a noise detection algorithm developed by Mallat and coworkers in signal analysis.

The resulting artificial viscosity can be incorporated into the difference scheme such that an entropy inequality holds.

PARALLEL PRECONDITIONER FOR CFD

Wei-Pai Tang

An effective linear solver is an essential part of sophisticated CFD code. Commonly, most of the CPU time of the computation is spent on the linear solver. When the size of the model grows and the difficulty of the simulation increases, the performance of the linear solver becomes crucial. If the computation is carried out in a high performance computer, new issues also arise.

The use of preconditioned Krylov space methods has been proven to be a competitive solution technique for wide ranges of large sparse matrix problems. It is acknowledged now that a high quality preconditioner is the key to the success. The objective of this project is to investigate parallel implementation issues when a high performance computer is used. This is a joint project with NAS Advanced Algorithm Division. It is an on-going project. First phase of prototype testing is completed. A real parallel implementation is planned for next year.

RESEARCH IN AERODYNAMIC SHAPE OPTIMIZATION

James Reuther and Antony Jameson

Since the inception of CFD, researchers have sought not only accurate aerodynamic prediction methods for given configurations, but also design methods capable of creating new optimum configurations. Yet, while flow analysis can now be carried out over quite complex configurations using the Navier-Stokes equations with a high degree of confidence, direct CFD based design is often limited to very simple two-dimensional and three-dimensional configurations, usually without the inclusion of viscous effects. The main effort of this research is to overcome the difficulties present in traditional aerodynamic optimization methods by introducing new technology. The CFD-based aerodynamic design methods of the past can be grouped into two basic categories: inverse methods, and numerical optimization methods.

Inverse methods derive their name from the fact that they invert the goal of the flow analysis algorithm. Instead of obtaining the surface distribution of an aerodynamic quantity, such as pressure, for a given shape, they calculate the shape for a given surface distribution of an aerodynamic quantity. Most of these methods are based on potential flow techniques, and few of them have been extended to three-dimensions. The common trait of all inverse methods is their computational efficiency. Typically, transonic inverse methods require the equivalent of 2-10 complete flow solutions in order to render a complete design. Since obtaining a few solutions for simple two-dimensional and three-dimensional designs can be done in at most a few hours on modern computers systems, the computational cost of most inverse methods is considered to be minimal. Unfortunately, they suffer from many limitations and difficulties. Their most glaring limitation is that the objective is built directly into the design process and thus cannot be changed to an arbitrary or more appropriate objective function.

A traditional alternative, which avoids some of the difficulties of inverse methods, but only at the price of heavy computational expense, is to use numerical optimization methods. The

essence of these methods is very simple: a numerical optimization procedure is coupled directly to an existing CFD analysis algorithm. The numerical optimization procedure attempts to extremize a chosen aerodynamic measure of merit which is evaluated by the chosen CFD code. Most of these optimization procedures require gradient information in addition to evaluations of the objective function. Here, the gradient refers to changes in the objective function with respect to changes in the design variables. The simplest method of obtaining gradient information is by finite differences. In this technique, the gradient components are estimated by independently perturbing each design variable with a finite step, calculating the corresponding value of the objective function using CFD analysis, and forming the ratio of the differences. These methods are very versatile, allowing any reasonable aerodynamic quantity to be used as the objective function. They can be used to mimic an inverse method by minimizing the difference between target and actual pressure distributions, or may instead be used to maximize other aerodynamic quantities of merit such as L/D. Unfortunately, these finite difference numerical optimization methods, unlike the inverse methods, are computationally expensive because of the large number of flow solutions needed to determine the gradient information for a useful number of design variables. Tens of thousands of flow analyses would be required for a complete three-dimensional design.

In our research, a new method is developed that avoids the limitations and difficulties of traditional inverse methods while retaining their inherent computational efficiency. The method dramatically reduces the cost of aerodynamic optimization by replacing the expensive finite-difference method of calculating the required gradients with an adjoint variable formulation. After deriving the differential form of the adjoint equations and posing the correct boundary conditions based on the objective function, the resulting system is discretized and solved on the same mesh as that used for the flow solution. A significant economization is thus achieved by applying the same subroutines used for the flow solution to the solution of the adjoint equations. The resulting design process requires only one flow calculation and one adjoint calculation per gradient evaluation, as opposed to the hundreds required for a finite-difference gradient involving hundreds of design variables. In practice the computational cost of the new method is two orders of magnitude less than a conventional approach.

Considerable effort has been focused in the last two years to develop control theory-based aerodynamic shape optimization methods. The work that has taken place in the last two years can be broken down into three specific areas.

- 1) Two-dimensional proof-of-concept studies.
- 2) Three-dimensional demonstration and research tool development.
- 3) NASA and industrial evaluation and feedback.

During the first year, work was primarily focused in area (1) and to a lesser extent area (2). At the beginning of this program at RIACS, methods were in place which showed that control theory could be used in conjunction with numerical optimization and computational fluid dynamics to create efficient design tools for flows governed by the potential flow equation (AIAA Paper 94-0499).

During the course of the first year of the program the development of adjoint methods was extended to treat the Euler equations. In our paper at the Multi-Disciplinary Optimization conference during summer 1994 (AIAA paper 94-4272, also RIACS report 94.18) results

were shown that demonstrated that control theory could be used to design airfoils that operate under transonic conditions by employing both an analytic mapping and a general mesh perturbation approach. Various objective functions were demonstrated showing the versatility of the new method. In the work presented at VKI, the first examples of three-dimensional wing design using control theory were presented. Finally, in a paper presented at the January 1995 Aerospace Sciences Meeting (AIAA paper 95-0123, also RIACS report 95.01) results for the design of wing and wing-body configurations over general meshes were shown.

During the last year (second year of the program) work continued in area (2) but strong emphasis was placed in area (3). One of the successes of the area (3) effort involved the participation of Beechcraft Aircraft Division of Raytheon, Inc. Raytheon entered into a cooperative agreement with NASA Ames Research Center to explore the usefulness of the adjoint-based design optimization methods we have been developing. In the middle of March 1995, representatives from Raytheon were on-site at Ames to test the adjoint based design techniques on a new transonic wing that they were developing for an all-new business jet they proposed to build and market. Since at that time we were able to treat only the design of wing-body configurations and their design involved a wing-body-nacelle configuration, some very imaginative design strategies were developed in order to permit our methods to be applicable to their very complex real world problem. The main goal of the cooperative effort was to test the adjoint design methods for the case of wing design. However, since the nacelles on their new configuration were very closely located to the wing, they proved to have a very strong influence on the flow solution about the wing at transonic speeds. Thus the design process had to incorporate the effect that these tightly coupled nacelles had on the wing solution. The remedy used by the Ames/Raytheon design team was to model the existence of the nacelles by a bump placed on the side of the fuselage body. The shape and extent of this bump was itself designed by the use of the adjoint method where a target pressure optimization was carried out using targets obtained from another CFD code which analyzes the complete configuration including the nacelles. Once an appropriate bump had been designed which mimicked the presence of the nacelles for the baseline wing, the adjoint-based design method was turned loose to reshape the wing with the fuselage including the body bump, which remained fixed. Several iterations with slight bump modifications and wing redesigns were carried out in a very rapid design cycle effort. By the beginning of May, a new wing had been designed using the new technology and validated computationally using another CFD code. This one-month design of a new transonic wing compares with the usual development time of more than a year for traditional methods. Raytheon has since wind tunnel tested the new wing design to confirm its predicted performance, and launched the design for production. They took 51 orders for the new airplane on the day they announced the design. Furthermore, Raytheon has been so impressed by the capability of adjoint-based design methods that they are now incorporating them into their own aircraft design environments. A paper authored by both NASA and Raytheon personnel that presents the basic design strategy and its outcome was presented at the Aerospace Sciences Meeting, January 1996 (AIAA paper 96-0554, also RIACS 96.03).

Another group that has taken a keen interest in our research is the NASA High Speed Research Program (HSR) group. In their effort to create economically viable supersonic transport configurations for the next century they are investigating the use of aerodynamic shape optimization to improve aerodynamic performance. Both the traditional as well as adjoint-based design methods created by our group at Ames have been tested by the HSR community. During the summer of 1995 our group supported many wind tunnel tests at NASA Langley Research Center in Virginia to validate the aero-performance gains

predicted by our aerodynamic shape optimization capability. A paper that gives an example of the capabilities of this emerging technology for supersonic design was presented at the American Society of Mechanical Engineers annual winter meeting in November 1995 (also RIACS 95.14).

The experience of working with Raytheon and the HSR community has forced us to consider many aspects of aerodynamic shape design that are often neglected from the purely academic stand-point. One of the many new areas of research that will develop as adjoint-based design becomes established is the study of possible aerodynamic design space parameterizations. We have made a first look into this area through a paper that explores the use of both Hicks-Henne functions and B-spline control points as design variables. The paper was presented at the Sixth International Symposium on Computational Fluid Dynamics conference in September 1995 (also RIACS report 95.13). In another effort to enhance our capabilities, the methodology was extended for the treatment of multiblock structured three-dimensional meshes. To accomplish this task we devised three new elements: a multiblock flow solver, a multiblock adjoint solver, and a multiblock mesh perturbation method. This new implementation will allow the design of complete aircraft configurations to be treated as part of the design process. This method can now treat the complex wing-body-nacelle design needed by Raytheon without resorting to difficult fictitious nacelle strategies. A paper was presented at the 34th Aerospace Sciences Meeting (AIAA paper 96-0094, also RIACS report 96.02) which presents this new multiblock design method.

In spite of the major accomplishments achieved during the last two years much more research will be required in order to harness the true potential of adjoint-based aerodynamic design. Since the time of our last paper in January 1996, work has been on-going to port this developing technology to parallel computer platforms and thus further reduce the design cycle time. A paper was presented in beginning of September at the Multi-Disciplinary Optimization Conference (AIAA paper 96-4045) that demonstrated the combined efficiency of parallel computing and adjoint methods for the aerodynamic design of a complete high speed transport configuration. This represents the first such calculation ever performed and involved the simultaneous design of the wing surface shape and the nacelle/diverter integration. This complete configuration design tool, which use both efficient parallel computations and the adjoint method, will make its way in the next year from the research level that it is currently in, to the production level that is necessary for the HSR community.

Further research is also underway in using the three-dimensional single-block implementation as a test bed, developments are being studied which will allow both constrained and multipoint formulations to be incorporated within the framework of adjoint-based design. Concurrently, investigations are under way to extend the development of the adjoint to treat the Navier-Stokes equations. Even with much of this work still to be accomplished it is nevertheless gratifying that the developments that have been achieved thus far have demonstrated beyond a doubt the great value of adjoint-based aerodynamic design. It is hoped that with all of these advances, the greater aeronautical science community will in the future adopt these new ideas into their production design environments. Certainly if the work in conjunction with Raytheon is any indication, this is already taking place.

ADAPTIVE REFINEMENT OF COMPOSITE CURVILINEAR GRIDS

Steven Suhr

A software system is being designed and implemented to manage the adaptive refinement of composite curvilinear grids for the approximate solution of time-dependent partial differential equations. With the simplifying assumption that the spatial domain has fixed boundaries, an initial grid is constructed using a fixed set of overlapping grids, which collectively conform to the boundaries and cover the domain. Refinement grids, aligned with the original base grids, are added to maintain accuracy as the solution evolves. This approach organizes the grids into a geometrically nested tree of connected components, and it explores the use of curvilinear stairstep refinement grids.

The programming language Vorpal, currently under development, will be used in the implementation of this system, taking advantage of Vorpal's support for data structures, abstract data types, structured external files, and modular program structure. An important milestone in the transformation of Vorpal from a collection of useful concepts into a unified preliminary design will be the translation of adaptive grid code being implemented at Stanford by others for model problems in two space dimensions, from C into Vorpal. When an implementation of Vorpal exists, the adaptive grid system will be extended from two to three space dimensions, and a version which can be applied more readily to realistic and diverse problems will be created. As the adaptive grid system evolves and grows, the anticipated future support in Vorpal for concurrency should also be useful.

CARTESIAN GRID METHODS FOR COMPLEX GEOMETRY

Marsha Berger

We are developing algorithms to simulate steady state flows in three space dimensions using a Cartesian grid representation of the geometry. This is in collaboration with Captain Michael Aftosmis, at Wright-Patterson/NASA Ames Research Center, and John Melton, at NASA Ames Research Center. In this approach, a solid object is superimposed on an underlying Cartesian grid, and the flow is computed around the object. This makes the problem of volume grid generation substantially easier, with the bulk of the work reduced to finding intersections between a possibly complex configuration and a regular Cartesian grid. However, the difficulty of grid generation is traded for the difficulties in the flow solver of imposing solid wall boundary conditions on a non-body fitted grid. Our previous work on flow solvers for this kind of grid however indicates that acceptable results that maintain second order accuracy over the entire flow field can be obtained.

Our research this summer is focusing on two important issues still needing attention. The first concerns robustness in the computational geometry procedures which form the basis of the grid generation algorithms. We are investigating the use of high precision and integer arithmetic packages for those cases when forward error analysis indicates an untrustworthy computation. This will be coupled with an algorithm to handle degeneracies, i.e. those cases with borderline results that take 99% of the programming but rarely occur. The second issue is the inherent efficiency of a Cartesian grid Euler solver. Previous studies have shown that Cartesian non-body-fitted grids typically take on the order of 20-25% more grid points to achieve the same accuracy as a body-fitted grid. These studies included the use of isotropically adaptively refined cells to resolve the geometry and the flow solution. We are investigating the use of directional adaptation to improve the efficiency, along with redesigning the data structures to reduce memory requirements.

DYNAMIC SPECTRAL PARTITIONING FOR UNSTRUCTURED ADAPTIVE GRID COMPUTATIONS

Andrew Sohn

The computational requirements for an adaptive solution of unsteady problems change as the simulation progresses. This change causes workload imbalance among processors on a parallel machine which, in turn, requires significant data movement at runtime. I have been developing a dynamic load-balancing framework, called JOVE, that balances the workload across all processors with a global view.

One of the key modules in JOVE is a mesh partitioner which partitions a computational mesh into pieces to measure the computational load of each processor. While I was visiting RIACS, I have been developing a new mesh partitioning method, called Dynamic Spectral Partitioning. This dynamic spectral bisection algorithm is based on the center of inertia of the unpartitioned dual graph vertices and utilizes information from the initial spectral partitioning. It is thus capable of rapidly updating a partition from one grid to the next to allow fast runtime load balancing.

A preliminary version of DSP has been developed and implemented on an IBM SP-2 distributed-memory machine. Some preliminary experimental results indicate that the dynamic spectral partitioner is twice faster than the best performing partitioner while yielding comparable solution quality.

NUMERICAL METHODS FOR THE COMPRESSIBLE NAVIER-STOKES EQUATIONS WITH APPLICATIONS TO AERODYNAMIC FLOWS

David Zingg

David Zingg continued his collaborative work with Tom Pulliam of the NASA Ames Research Center on numerical methods for the compressible Navier-Stokes equations, with applications to aerodynamic flows. Topics studied include local preconditioning, matrix dissipation, Newton-Krylov methods, and multigrid methods. Dr. Zingg gave a seminar at Ames entitled "Research in CFD and Related Interdisciplinary Areas at the University of Toronto Institute for Aerospace Studies," which covered algorithm development for computational aerodynamics, turbulence model studies for aerodynamic flows at high lift, and numerical methods for simulating wave phenomena applied to the time-domain Maxwell equations.

B. HIGH PERFORMANCE NETWORKS

BAY AREA GIGABIT NETWORK TESTBED

Marjory J. Johnson

The planning, development and use of the Bay Area Gigabit Network Testbed (BAGNet) has been a major project during this contract period. M. Johnson, Bill Johnston of Lawrence Berkeley Laboratory, and Dan Swinehart of Xerox PARC are co-coordinators of the BAGNet project.

An effort to establish a gigabit testbed within the Bay Area began in the late 1980s. Fourteen organizations within the Bay Area have been involved, including major technology companies, research organizations, universities, and government laboratories. M. Johnson has collaborated with persons from the Numerical Aerodynamic Simulation Systems Division (NAS) and the NASA Science Internet Project Office (NSIPO) to represent NASA in this effort. M. Johnson played a key role in writing many of the early proposals and descriptive documents that were required during the testbed-planning phase.

After investigating several alternatives for funding the testbed infrastructure, an opportunity was presented through Pacific Bell's CalREN (California Research and Education) program. This program was initiated in 1993 for the purpose of stimulating the development and dissemination of high-speed communications applications within the state of California. CalREN-sponsored testbeds would be provided access to Pacific Bell's ATM network services, a necessity for our metropolitan-area testbed.

BAGNet was the first of the CalREN-sponsored testbeds to be implemented. Physical deployment of the testbed began on December 30, 1993, when the first two testbed sites, NASA Ames and Xerox Palo Alto Research Center (PARC), were connected to the Pacific Bell ATM switch via OC-3c links. The remaining thirteen sites were gradually added to the testbed during 1994. Each site is connected by an OC-3c (155 Mbps) link to the Pacific Bell infrastructure. Up to four hosts per site are directly connected to the testbed. A mesh of permanent virtual channels provides connections between all possible pairs of hosts. BAGNet is an IP over ATM testbed, with AAL5 as the adaptation layer.

The featured application for BAGNet is the teleseminar application. Our initial goal was to implement simple multicast video transmissions. We set up point-to-multipoint permanent virtual channels, to enable each host on BAGNet to broadcast to all BAGNet sites. This enables us to simulate multicast until a better solution is developed within the standards groups. Experimental multicast over these point-to-multipoint virtual channels began in 1994.

We have used the World Wide Web (WWW) to coordinate testbed activities and to post information about BAGNet. The URL for the general BAGNet home page on WWW, which is maintained by Lawrence Berkeley Laboratory (LBL), is <http://george.lbl.gov/BAGNet.html>. Site-specific BAGNet home pages, which contain information about the individual testbed sites, can be accessed from the LBL home page. The RIACS testbed URL is ftp://riacs.edu/pub/Gigabit_Testbed/psyche-ping.html.

BAGNet has been an ambitious project, both because of the immature level of ATM technology when the testbed originated, and because of the large number of testbed

participants and the heterogeneity of equipment involved. Major testbed accomplishments during 1995 included identification and resolution of many issues that must be addressed when building a relatively large-scale IP-over-ATM network, development of a methodology to maintain up-to-date configuration status, analysis of performance degradation caused by bandwidth-policing policies, compilation of comprehensive performance statistics obtained by using various performance tools (e.g., Netperf), and implementation of multicast capabilities by using point-to-multipoint PVCs. Although we originally planned to use switched virtual channels as soon as possible, the technology is not yet ready. Hence, we have not been able to experiment with specifying and delivering quality-of-service guarantees. This has been a major disappointment within the testbed community.

During 1995 several testbed sites multicasted seminars on a regular basis over BAGNet. The quality of reception varied, depending on individual workstation capabilities. Clearly, the bottleneck for performance is not the network, but rather workstation architecture and protocol issues.

Late in 1995 we initiated a collaboration with Bellcore to capture data for traffic analysis, so that we can see what a data stream for a real ATM application looks like. Bellcore is interested in general traffic analysis to help them understand how to manage data transmission for commercial-service offerings; several of the individual sites plan to use the collected data to model specific applications. We are also interested in observing and analyzing interactions between applications that are sharing the testbed.

One measurement period was conducted in September. Due to problems with the measurement equipment, a second measurement period was conducted in April. M. Johnson participated in both these sessions. We are currently analyzing this data.

Early in 1996 M. Johnson upgraded her BAGNet workstation from a Sun SparcStation 2 clone to a Sun SparcStation 20 clone. This process was frustrating and time-consuming, due to the difference in CPU speed of the clone and a real Sun SparcStation 20. The associated problems have not yet been completely resolved.

At RIACS, in addition to participating in testbed-wide activities, we are developing and testing protocols for file transfer that will utilize the available high bandwidth of the testbed. We are also developing an environment to support collaborative scientific work. Other testbed projects within NASA Ames include video on demand (VOD) and remote access to NASA wind tunnels.

BAGNet participants have both generated and responded to worldwide interest in this project by contributing numerous publications and presentations to the networking community. M. Johnson has written the following papers:

“Experiences with the Bay Area Gigabit Network Testbed,”
 Proceedings of the 5th IEEE Workshop on Future Trends in
 Distributed Computer Systems, Aug. 1995, pp. 26-32.

“Achieving High Throughput for Data Transfer over ATM
 Networks,” Proceedings of the International Communications
 Conference, ICC'96, June, 1996 (with Jeffrey N. Townsend).

COLLABORATIVE SCIENCE

Marjory J. Johnson

The objective of this project, initiated in 1994, is to develop and analyze new collaborative working paradigms for earth and space scientists by combining data-analysis tools with state-of-the-art networking tools. These new paradigms will address the use of massive data sets that are stored remotely, collaboration between geographically separated colleagues, and joint data analysis.

Project collaborators have included earth scientists at NASA Ames and at the University of Arizona, networking researchers at Sandia National Laboratories - Livermore, and space scientists at Lockheed Palo Alto Research Laboratory (PARL). Several RIACS student employees have made key contributions to this project.

During 1994 we established the project infrastructure. We developed a collaborative work environment using a workstation which is directly connected to BAGNet (called the BAGNet workstation in the remainder of this text). We incorporated both a commercial data-analysis tool and a proprietary one, but have found them both to be unsuitable for collaborative use. We have incorporated software tools developed by Sandia to enable audio/video teleconferencing and the sharing of X-window applications.

During experiments over BAGNet between RIACS and Sandia in 1995, it was immediately apparent that the Sparc 2 architecture of our BAGNet workstation was inadequate to take full advantage of the bandwidth capabilities offered by BAGNet. This motivated the upgrade of our workstation to a Sparc 20 architecture.

In 1996 we developed an image-browsing tool for viewing large images. The standard paradigm used in conventional browsers is to download the entire image in scratch disk space provided at the client's site, and then to view the image by interacting with the local store. It is not feasible to use such browsers for very large image files, both because of the long time required to download the image and because of limitations on client disk space.

Our tool is designed specifically to address these issues. It allows large images to be viewed in real time, without requiring any scratch space to be provided by the client. In a typical session using this tool a user browses through images in the directory of a remote server, perhaps searching for an image of a particular geographical location or an image from a specific instrument. Upon locating a suitable candidate, the user selects this image for display on his workstation. As the user scrolls to various sections of the image, data for the display is transferred over the network from the server. A rectangle on an iconified version of the image identifies the portion of the image that is currently displayed.

We are experimenting with this tool over BAGNET. The tool is specially suited for accessing images over high-speed networks, since it is often considerably faster to access data from a remote system over a high-speed network than to retrieve it from a local network file server.

As part of the recent Bellcore traffic-measurement experiments we conducted typical collaborative work sessions, using our image-browsing tool and using Sandia's video-conferencing tools. We are currently analyzing this traffic data, to obtain an understanding of communication patterns for collaborative applications.

HIGH-RATE DATA TRANSFER

Marjory J. Johnson

The standard file-transfer protocol, ftp, yields throughput rates that are disappointingly low, relative to the raw bandwidth that is available with high-performance fiber-optic networks such as the BAGNet testbed. Disk I/O is clearly a primary bottleneck, and sender/receiver interactions to ensure error-free transmission provide another source of significant overhead. Such issues indicate that new data-transfer protocols must be designed to utilize the bandwidth that is available with emerging network technologies.

Our application focus is the transfer of large image files. Since many image-transfer applications can tolerate a low level of transmission errors, we are basing our protocol development on UDP rather than TCP. Our goal is to develop a data-transfer protocol that maximizes network throughput over ATM networks, while keeping transmission errors manageable. Of course, the level of transmission errors that is considered acceptable is application dependent.

We are currently experimenting with several techniques for data transfer, all of which attempt to keep the transmission pipe full. We are using multiple data streams, so that disk I/O, etc. can be overlapped with data transfer. Transmission errors are controlled via a low level of synchronization of sender/receiver activities.

We are conducting our experiments on a variety of workstations located at three BAGNet sites: NASA Ames, Sandia National Laboratories, and Lawrence Livermore National Laboratory. Preliminary results are contained in the paper, "Achieving High Throughput for Data Transfer over ATM Networks," which was presented at the International Communications Conference, ICC '96.

Early results validate our approach. We are able to limit transmission errors to two to three percent, while achieving throughput rates that are several times higher than ftp rates.

MISCELLANEOUS PROJECTS

M. Johnson was co-investigator on a RIACS/University of Arizona collaboration on "Content-based Query and Browse of Earth Science Imagery Databases using High Performance Computers and Networks," a project funded by NASA under the HPCC/ESS (Earth System Science) Program. This project has led to an ongoing collaboration between RIACS and the University of Arizona.

M. Johnson has directed the work of numerous summer students. Justin Paola, University of Arizona, investigated the use of neural networks for the classification of remotely sensed multispectral imagery. He examined the use of parallel algorithms on various NASA computers. Michael Kumbera, University of Wisconsin - Milwaukee, implemented the NAS multigrid benchmark on the CM-5. His accomplishments provided a foundation for future work to analyze the performance of the interprocessor communication architecture of the CM-5 and other parallel computers. Marc Bumble, Pennsylvania State University, assisted in the collaborative science project. Jeffrey Townsend, Stanford University, has participated in several projects involving BAGNet. Jason Deng, who will begin his studies at Stanford University this fall, is currently assisting in the analysis of BAGNet traffic data.

M. Johnson has participated in a joint NASA/DoD effort to develop interoperable data communications standards for use in both civil and military space projects. She is currently a member of the ISO/TC20/USSCAG13 committee to develop communication standards for space missions.

M. Johnson has been active in the general networking community, by helping to organize conferences, serving on various program committees, and by refereeing papers. She also served on a review committee for a DoE collaboratory program, and a review panel for NSF.

III. TECHNICAL REPORTS

96.01 A HIGH-ORDER FINITE DIFFERENCE SCHEME WITH SHARP SHOCK RESOLUTION FOR THE EULER EQUATIONS

Margot Gerritsen and Pelle Olsson (Stanford University)
January 1996 (27 pages)

We derive a high-order finite difference scheme for the Euler equations that satisfies a semi-discrete energy estimate, and present an efficient strategy for the treatment of discontinuities that leads to sharp shock resolution. The formulation of the semi-discrete energy estimate is based on a symmetrization of the Euler equations that preserves the homogeneity of the flux vector, a canonical splitting of the flux derivative vector, and the use of difference operators that satisfy a discrete analogue to the integration-by-parts procedure used in the continuous energy estimate. Around discontinuities or sharp gradients, refined grids are created on which the discrete equations are solved after adding a newly constructed artificial viscosity. The positioning of the sub-grids and computation of the viscosity are aided by a detection algorithm which is based on a multi-scale wavelet analysis of the pressure grid function. The wavelet theory provides easy to implement mathematical criteria to detect discontinuities, sharp gradients and spurious oscillations quickly and efficiently.

96.02 AERODYNAMIC SHAPE OPTIMIZATION OF COMPLEX AIRCRAFT CONFIGURATIONS VIA AN ADJOINT FORMULATION

James Reuther, Antony Jameson (Princeton University), J. Farmer (Brigham Young University), L. Martinelli (Princeton University) and D. Saunders (Sterling Software)
January 1996 (17 pages)

This work describes the implementation of optimization techniques based on control theory for complex aircraft configurations. Here control theory is employed to derive the adjoint differential equations, the solution of which allows for a drastic reduction in computational costs over previous design methods. In our earlier studies it was shown that this method could be used to devise effective optimization procedures for airfoils, wings and wing-bodies subject to either analytic or arbitrary meshes. Design formulations for both potential flows and flows governed by the Euler equations have been demonstrated, showing that such methods can be devised for various governing equations. In our most recent works the method was extended to treat wing-body configurations with a large number of mesh points, verifying that significant computational savings can be gained for practical design problems. In this paper the method is extended for the Euler equations to treat complete aircraft configurations via a new multiblock implementation. New elements include a multiblock-multigrid flow solver, a multiblock-multigrid adjoint solver, and a multiblock mesh perturbation scheme. Two design examples are presented in which the new method is used for the wing redesign of a transonic business jet.

96.03 BUSINESS JET WING DESIGN USING AERODYNAMIC SHAPE OPTIMIZATION

James Reuther, John W. Gallman (NASA Ames Research Center), Neal Pfeiffer, William Forest and David Bernstorf (Raytheon Aircraft Company)

January 1996 (12 pages)

A new method that relies on computational fluid dynamics (CFD) and numerical optimization is used to design a transonic business jet wing. The first step of this new design method is to develop target pressures for a three-dimensional wing design using a two-dimensional airfoil optimization code (MSESLINDOP). This airfoil optimization method is fast enough to solve a six-point design problem that is representative of an entire aircraft mission in a few minutes. A full-potential finite element code with a solution adaptive Cartesian grid (TRANAIR) is used to analyze the wing-body-nacelle configuration and establish the influence of the fuselage mounted nacelles on the wing pressures. The blockage in the flow caused by these nacelles is approximated in a wingbody Euler CFD code (SYN87) with a large bump on the aft fuselage. The SYN87 code also solves an adjoint set of equations to evaluate the flowfield. These flowfield sensitivities enable three dimensional shape optimization in this study with a quasi-Newton optimization routine. The objective function used to design both the fuselage bum and the wing contours was a sum-of-squares of the difference between computed and target wing pressures. Finally, the surface contours are modified slightly with a computer aided drawing machine to reduce manufacturing complexity. Wind tunnel data from the Boeing Transonic Wind Tunnel is in very good agreement with the pressure distributions developed for the 20 ° swept wing considered in this study. This data shows that the design goals of natural laminar flow at a Mach number of 0.75 and minimum wave drag at a Mach number of 0.80 have been met and provides a validation of the design method developed in this study.

96.04 EFFICIENT HELICOPTER AERODYNAMIC AND AEROACOUSTIC PREDICTIONS ON PARALLEL COMPUTERS

Andrew M. Wissink (University of Minnesota), Anastasios S. Lyrantzis (Purdue University), Roger C. Strawn (US Army AFDD), Leonid Oliker and Rupak Biswas

January 1996 (14 pages)

This paper presents parallel implementations of two codes used in a combined CFD/Kirchhoff methodology to predict the aerodynamics and aeroacoustics properties of helicopters. The rotorcraft Navier-Stokes code, TURNS, computes the aerodynamic flowfield near the helicopter blades and the Kirchhoff acoustics code computes the noise in the far field, using the TURNS solution as input. The overall parallel strategy adds MPI message passing calls to the existing serial codes to allow for communication between processors. As a result, the total code modifications required for parallel execution are relatively small. The biggest bottleneck in running the TURNS code in parallel comes from the LU-SGS algorithm that solves the implicit system of equations. We use a new hybrid domain decomposition implementation of LU-SGS to obtain good parallel performance on the SP-2. TURNS demonstrates excellent parallel speedups for quasi-steady and unsteady three-dimensional calculations of a helicopter blade in forward flight. The execution rate attained by the code on 114 processors is six times faster than the same cases run on one processor of the Cray C-90. The parallel Kirchhoff code also shows excellent parallel speedups and fast execution rates. As a performance demonstration, unsteady acoustic pressures are computed at 1886 far-field observer locations for a sample acoustics problem. The calculation requires over two hundred hours of CPU time on one C-90 processor but takes only a few hours on 80 processors of the SP2. The resultant far-field acoustic field is

analyzed with state-of-the-art audio and video rendering of the propagating acoustic signals.

96.05 NUMERICAL CONFORMAL MAPPING USING CROSS-RATIOS AND DELAUNAY TRIANGULATION

Tobin A. Driscoll (Cornell University) and Stephen A. Vavasis (Cornell University)
January 1996 (32 pages)

We propose a new algorithm for computing the Riemann mapping of the unit disk to a polygon, also known as the Schwarz-Christoffel transformation. The new algorithm, CRDT, is based on cross-ratios of the prevertices, and also on cross-ratios of quadrilaterals in a Delaunay triangulation of the polygon. The CRDT algorithm produces an accurate representation of the Riemann mapping even in the presence of arbitrary long, thin regions in the polygon, unlike any previous conformal mapping algorithm. We believe that CRDT can never fail to converge to the correct Riemann mapping, but the correctness and convergence proof depend on conjectures that we have so far not been able to prove. We demonstrate convergence with computational experiments. The Riemann mapping has applications to problems in two-dimensional potential theory and to finite-difference mesh generation. We use CRDT to produce a mapping and solve a boundary value problem on long, thin regions for which no other algorithm can solve these problems.

96.06 AN OVERSET GRID NAVIER-STOKES/KIRCHHOFF-SURFACE METHOD FOR ROTORCRAFT AERACOUSTIC PREDICTIONS

Earl P. N. Duque (US Army AFDD), Roger C. Strawn (US Army AFDD), Jasim Ahmad (Sterling Software) and Rupak Biswas
January 1996 (13 pages)

This paper describes a new method for computing the flowfield and acoustic signature of arbitrary rotors in forward flight. The overall scheme uses a finite-difference Navier-Stokes solver to compute the aerodynamic flowfield near the rotor blades. The equations are solved on a system of overset grids that allow for prescribed cyclic and flapping blade motions and capture the interactions between the rotor blades and wake. The far-field noise is computed with a Kirchhoff integration over a surface that completely encloses the rotor blades. Flowfield data are interpolated onto this Kirchhoff surface using the same overset-grid techniques that are used for the flowfield solution. As a demonstration of the overall prediction scheme, computed results for far-field noise are compared with experimental data for both high-speed impulsive (HSI) and blade-vortex interaction (BVI) cases. The HSI case showed good agreement with experimental data while a preliminary attempt at the BVI case did not. The computations clearly show that temporal accuracy, spatial accuracy and grid resolution in the Navier-Stokes solver play key roles in the overall accuracy of the predicted noise. These findings will be addressed more closely in future BVI computations. Overall, the overset-grid CFD scheme provides a powerful new framework for the prediction of helicopter noise.

96.07 SATISFIABILITY TEST WITH SYNCHRONOUS SIMULATED ANNEALING ON THE FUJITSU AP1000 MASSIVELY-PARALLEL MULTIPROCESSOR

Andrew Sohn (NJIT) and Rupak Biswas
March 1996 (8 Pages)

Solving the hard Satisfiability Problem is time consuming even for modest-sized problem instances. Solving the Random L-SAT Problem is especially difficult due to the ratio of

clauses to variables. This report presents a parallel synchronous simulated annealing method for solving the Random L-SAT Problem on a large-scale distributed-memory multiprocessor. In particular, we use a parallel synchronous simulated annealing procedure, called Generalized Speculative Computation, which guarantees the same decision sequence as sequential simulated annealing. To demonstrate the performance of the parallel method, we have selected problem instances varying in size from 100-variables/425-clauses to 5000-variables/21,250-clauses. Experimental results on the AP1000 multiprocessor indicate that our approach can satisfy 99.9% of the clauses while giving almost a 70-fold speedup on 500 processors.

96.08 ACHIEVING HIGH THROUGHPUT FOR DATA TRANSFER OVER ATM NETWORKS

Marjory J. Johnson and Jeffrey N. Townsend

March 1996 (8 pages)

Proceedings of IEEE ICC '96, June 1996, pp. 405 - 411

File-transfer rates for ftp are often reported to be relatively slow, compared to the raw bandwidth available in emerging gigabit networks. While a major bottleneck is disk I/O, protocol issues impact performance as well. Ftp was developed and optimized for use over the TCP/IP protocol stack of the Internet. However, TCP has been shown to run inefficiently over ATM. In an effort to maximize network throughput, data-transfer protocols can be developed to run over UDP or directly over IP, rather than over TCP. If error-free transmission is required, techniques for achieving reliable transmission can be included as part of the transfer protocol. However, selected image-processing applications can tolerate a low level of errors in images that are transmitted over a network. In this paper we report on experimental work to develop a high-throughput protocol for unreliable data transfer over ATM networks. We attempt to maximize throughput by keeping the communications pipe full, but still keep packet loss under five percent. We use the Bay Area Gigabit Network Testbed as our experimental platform.

96.09 AERODYNAMIC SHAPE OPTIMIZATION USING CONTROL THEORY

James John Reuther

May 1996 (226 pages)

Aerodynamic shape design has long persisted as a difficult scientific challenge due to its highly nonlinear flow physics and daunting geometric complexity. However, with the emergence of Computational Fluid Dynamics (CFD) it has become possible to make accurate predictions of flows which are not dominated by viscous effects. It is thus worthwhile to explore the extension of CFD methods for flow analysis to the treatment of aerodynamic shape design. Two new aerodynamic shape design methods are developed which combine existing CFD technology, optimal control theory, and numerical optimization techniques. Flow analysis methods for the potential flow equation and the Euler equations form the basis of the two respective design methods. In each case, optimal control theory is used to derive the adjoint differential equations, the solution of which provides the necessary gradient information to a numerical optimization method much more efficiently than by conventional finite differencing. Each technique uses a quasi-Newton numerical optimization algorithm to drive an aerodynamic objective function toward a minimum. An analytic grid perturbation method is developed to modify body fitted meshes to accommodate shape changes during the design process. Both Hicks-Henne perturbation functions and B-spline control points are explored as suitable design variables. The new methods prove to be computationally efficient and robust, and can be used for practical airfoil design including geometric and aerodynamic constraints. Objective functions are

chosen to allow both inverse design to a target pressure distribution and wave drag minimization. Several design cases are presented for each method illustrating its practicality and efficiency. These include non-lifting and lifting airfoils operating at both subsonic and transonic conditions.

96.10 COMPUTATIONAL METHODS FOR THE PREDICTION AND ANALYSIS OF HELICOPTER NOISE

Roger C. Strawn (US Army AFDD), Leonid Oliker and Rupak Biswas
May 1996 (11 pages)

This paper describes several new methods to predict and analyze rotorcraft noise. These methods are: 1) a combined computational fluid dynamics and Kirchhoff scheme for far-field noise predictions, 2) parallel computer implementation of the Kirchhoff integrations, 3) audio and visual rendering of the computed acoustic predictions over large far-field regions, and 4) acoustic tracebacks to the Kirchhoff surface to pinpoint the sources of the rotor noise. The paper describes each method and presents sample results for three test cases. The first case consists of in-plane high-speed impulsive noise and the other two cases show idealized parallel and oblique blade-vortex interactions. The computed results show good agreement with available experimental data but convey much more information about the far-field noise propagation. When taken together, these new analysis methods exploit the power of new computer technologies and offer the potential to significantly improve our prediction and understanding of rotorcraft noise.

96.11 PARALLEL IMPLEMENTATION OF AN ADAPTIVE SCHEME FOR 3D UNSTRUCTURED GRIDS ON THE SP2

Leonid Oliker, Rupak Biswas and Roger C. Strawn (US Army AFDD)
May 1996 (13 pages)

Dynamic mesh adaption on unstructured grids is a powerful tool for computing unsteady flows that require local grid modifications to efficiently resolve solution features. For this work, we consider an edge-based adaption scheme that has shown good single-processor performance on the C90. We report on our experience parallelizing this code for the SP2. Results show a 47.0X speedup on 64 processors when 10% of the mesh is randomly refined. Performance deteriorates to 7.7X when the same number of edges are refined in a highly-localized region. This is because almost all the mesh adaption is confined to a single processor. However, this problem can be remedied by repartitioning the mesh immediately after targeting edges for refinement but before the actual adaption takes place. With this change, the speedup improves dramatically to 43.6X.

96.12 A REVIEW OF HIGH-ORDER AND OPTIMIZED FINITE-DIFFERENCE METHODS FOR SIMULATING LINEAR WAVE PHENOMENA

David W. Zingg (University of Toronto Institute for Aerospace Studies)
June 1996 (29 pages)
Submitted to the AIAA 13th CFD conference and Journal of Computational Physics.

This paper presents a review of high-order and optimized finite-difference methods for numerically simulating the propagation and scattering of linear waves, such as electromagnetic, acoustic, or elastic waves. The spatial operators reviewed include compact schemes, noncompact schemes, schemes on staggered grids, and schemes which are optimized to produce specific characteristics. The time-marching methods discussed include Runge-Kutta methods, Adams-Bashforth methods, and the leapfrog method. In addition, the following fourth-order fully-discrete finite-difference methods are considered:

a one-step implicit scheme with a three-point spatial stencil, a one-step explicit scheme with a five-point spatial stencil, and a two-step explicit scheme with a five-point spatial stencil. For each method studied, the number of grid points per wavelength required for accurate simulation of wave propagation over large distances is presented. Recommendations are made with respect to the suitability of the methods for specific problems and practical aspects of their use, such as appropriate Courant numbers and grid densities. Avenues for future research are suggested.

96.13 UNSTRUCTURED ADAPTIVE GRID COMPUTATIONS ON AN ARRAY OF SMPS

Rupak Biswas, Ira Pramanick (SGI) Andrew Sohn (NJIT) and Horst Simon (NERSC - LBL)

July 1996 (8 pages)

Proceedings of Parallel CFD '96, May 1996

Dynamic load balancing is necessary for the parallel adaptive solution of unsteady problems in fluid dynamics, since their computational requirements change as the simulation progresses leading to load imbalance. JOVE is such a dynamic load-balancing framework. We study the performance of two different implementations of JOVE on the Silicon Graphics' POWER CHALLENGE array. This parallel machine is an array of shared-memory symmetric multiprocessing (SMP) systems, an architecture that is becoming increasingly popular as the most useful model of scaleable parallel computing. Parallel algorithms need to be designed to exploit the hybrid communication model offered by such an architecture, and in this paper, we study these issues as they relate to JOVE.

96.14 ALGORITHMS FOR AUTOMATIC ALIGNMENT OF ARRAYS

Leonid Oliker, Siddhartha Chatteerjee (University of North Carolina), John R. Gilbert (Xerox PARC), Robert Schreiber (Hewlett-Packard Company) and Thomas J. Sheffler (Rambus, Inc.)

August 1996 (32 pages)

Appeared in a special issue of the Journal of Parallel and Distributed Computing, August 1996.

Aggregate data objects (such as arrays) are distributed across the processor memories when compiling a data-parallel language for a distributed-memory machine. The mapping determines the amount of communication needed to bring operands of parallel operations into alignment with each other. A common approach is to break the mapping into two stages: an alignment that maps all the objects to an abstract template, followed by a distribution that maps the template to the processors. This paper describes algorithms for solving the various facets of the alignment problem: axis and stride alignment, static and mobile offset alignment, and replication labeling. We show that optimal axis and stride alignment is NP-complete for general program graphs, and give a heuristic method that can explore the space of possible solutions in a number of ways. We show that some of these strategies can give better solutions than a simple greedy approach proposed earlier. We also show how local graph contractions can reduce the size of the problem significantly without changing the best solution. This allows more complex and effective heuristics to be used. We show how to model the static offset alignment problem using linear programming, and we show that loop-dependent mobile offset alignment is sometimes necessary for optimum performance. We describe an algorithm with for determining mobile alignments for objects within do loops. We also identify situations in which replicated alignment is either required by the program itself or can be used to improve performance. We describe an algorithm based on network flow that replicates objects so as

to minimize the total amount of broadcast communication in replication. We present experimental results showing the effect of our axis/stride alignment algorithm on the performance of some example programs running on the CM-5.

96.15 IMPACT OF LOAD BALANCING ON UNSTRUCTURED ADAPTIVE GRID COMPUTATIONS FOR DISTRIBUTED-MEMORY MULTIPROCESSORS

Rupak Biswas, Andrew Sohn (NJIT) and Horst Simon (NERSC - LBL)

July 1996 (8 pages)

To appear Proceedings of 8th IEEE Symposium on Parallel and Distributed Processing, October 23-16, 1996

The computational requirements for an adaptive solution of unsteady problems change as the simulation progresses. This causes workload imbalance among processors on a parallel machine which, in turn, requires significant data movement at runtime. We present a new dynamic load-balancing framework, called JOVE, that balances the workload across all processors with a global view. Whenever the computational mesh is adapted, JOVE is activated to eliminate the load imbalance. JOVE has been implemented on an IBM SP2 distributed-memory machine in MPI for portability. Experimental results for two model meshes demonstrate that mesh adaption with load balancing gives more than a sixfold improvement over one without load balancing. We also show that JOVE gives a 24-fold speedup on 64 processors compared to sequential execution.

96.16 GLOBAL LOAD BALANCING WITH PARALLEL MESH ADAPTION ON DISTRIBUTED-MEMORY SYSTEMS

Rupak Biswas, Leonid Oliker and Andrew Sohn (NJIT)

August 1996 (17 pages)

To appear Supercomputing '96, November 1996

Dynamic mesh adaption on unstructured grids is a powerful tool for efficiently computing unsteady problems to resolve solution features of interest. Unfortunately, this causes load imbalance among processors on a parallel machine. This paper describes the parallel implementation of a tetrahedral mesh adaption scheme and a new global load balancing method. A heuristic remapping algorithm is presented that assigns partitions to processors such that the redistribution cost is minimized. Results indicate that the parallel performance of the mesh adaption code depends on the nature of the adaption region and show a 35.5X speedup on 64 processors of an SP2 when 35% of the mesh is randomly adapted. For large-scale scientific computations, our non-balanced loads. Furthermore, our heuristic remapper yields processor assignments that are less than 3% off the optimal solutions but requires only 1% of the computational time.

96.17 AERODYNAMIC SHAPE OPTIMIZATION OF SUPER SONIC AIRCRAFT CONFIGURATIONS VIA AN ADJOINT FORMULATION ON DISTRIBUTED MEMORY PARALLEL COMPUTERS

James Reuther, Juan Jose Alonso (Princeton University), Mark J. Rimlinger (Simco) and A. Jameson (Princeton University)

September 1996 (18 pages)

This work describes the application of a control theory-based aerodynamic shape optimization method to the problem of supersonic aircraft design. The design process is greatly accelerated through the use of both control theory and a parallel implementation on distributed memory computers. Control theory is employed to derive the adjoint differential equations whose solution allows for the evaluation of design gradient

information at a fraction of the computational cost required by previous design methods. The resulting problem is then implemented on parallel distributed memory architectures using a domain decomposition approach, an optimized communication schedule, and the MPI (Message Passing Interface) Standard for portability and efficiency. The final result achieves very rapid aerodynamic design based on higher order computational fluid dynamics methods (CFD). In our earlier studies, the serial implementation of this design method was shown to be effective for the optimization of airfoils, wings, wing-bodies, and complex aircraft configurations using both the potential equation and the Euler equations. In our most recent paper, the Euler method was extended to treat complete aircraft configurations via a new multiblock implementation. Furthermore, during the same conference, we also presented preliminary results demonstrating that this basic methodology could be ported to distributed memory parallel computing architectures. In this paper, our concern will be to demonstrate that the combined power of these new technologies can be used routinely in an industrial design environment by applying it to the case study of the design of typical supersonic transport configurations. A particular difficulty of this test case is posed by the propulsion/airframe integration.

IV. PUBLICATIONS

Rupak Biswas, Leonid Oliker and Andrew Sohn (NJIT), "Parallel Mesh Adaption With Global Load Balancing On The SP2," Proceedings of the NASA Computational Sciences Workshop, page 76, August 1996.

Rupak Biswas, Leonid Oliker and Andrew Sohn (NJIT), "Global Load Balancing With Parallel Mesh Adaption On Distributed-Memory Systems," To appear in the Proceedings of Supercomputing 1996.

Susan E. Cliff (NASA Ames Research Center), Timothy J. Baker (Princeton University), Raymond M. Hicks (NASA Ames Research Center), and James Reuther, "Computational/Experimental Study Of Two Optimized Supersonic Transport Designs And The Reference H Baseline," NASA CDTM-21008, July 1996.

Marjory J. Johnson and Jeff N. Townsend, "Achieving High Throughput For Data Transfer Over ATM Networks," Proceedings of the 1996 IEEE International Conference on Communications, June 1996, pages 405-411.

Joseph Oliger and Xiaolei Zhu, "Stability And Error Estimation For Component Adaptive Grid Methods," Applied Numerical Mathematics volume 20 (1996), pages 407-426.

Leonid Oliker, Rupak Biswas and Roger Strawn (US Army AFDD), "Parallel Implementation Of An Adaptive Scheme For 3D Unstructured Grids On The SP2," Proceedings of Third International Workshop on Parallel Algorithms for Irregularly Structured Problems, pages 34-47, August 1996.

James Reuther, "Aerodynamic Shape Optimization Using Control Theory " Ph.D. Dissertation, University of California Davis, June 1996.

James Reuther, Juan J. Alonso, Antony Jameson (Princeton University), and Mark J. Rimlinger (Simco), "Rapid Cycle Aerodynamic Shape Optimization Of Complete Aircraft Configurations Via An Adjoint Formulation And Parallel Computing," Proceedings of the 1996 Computational Aerosciences Workshop, August 1996.

James Reuther, Juan J. Alonso, Antony Jameson (Princeton University), and Mark J. Rimlinger (Simco), "Aerodynamic Shape Optimization Of Supersonic Aircraft Configurations Via An Adjoint Formulation On Distributed Memory Parallel Computers," Proceedings of the 6th AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, September 1996 (AIAA paper 96-4045).

James Reuther, John W. Gallman (NASA Ames Research Center), Neal Pfeiffer, William Forest and David Bernstorff (Raytheon Aircraft Company), "Business Jet Wing Design

Using Aerodynamic Shape Optimization," Proceedings of The 34th Aerospace Sciences Meeting, January 1996 (AIAA paper 96-0554).

James Reuther, Antony Jameson (Princeton University), James Farmer (Brigham Young University), Luigi Martinelli (Princeton University) and David Saunders (Sterling Software), "Aerodynamic Shape Optimization Of Complex Aircraft Configurations Via An Adjoint Formulation," Proceedings of The 34th Aerospace Sciences Meeting, January 1996 (AIAA paper 96-0094).

James Reuther, David Saunders (Sterling Software), "Advances In Design Optimization Using Adjoint Methods," First NASA/Industry High-Speed Research Configuration Aerodynamics Workshop, February 1996.

Roger Strawn (US Army AFDD), Leonid Oliker and Rupak Biswas, "New Computational Methods For The Predication And Analysis Of Helicopter Noise," AIAA-96-1696, May 1996.

W.P. Tang, J. A. George (Univ. of Waterloo), K. Ikramov (Univ. of Moscow) and V.N. Tchugunov (Univ. of Moscow), "On Doubly Symmetric Tridiagonal Forms For Complex Matrices," SIAM Journal On Matrix Analysis And Applications, 1996 Vol. 17, No 3, pp. 680-690,

W.P. Tang and H. Sun (Univ. of Waterloo), "An Overdetermined Schwarz Alternating Method," SIAM Journal on Scientific Computing, Vol. 17, No. 4, pp. 884-905, 1996.

W.P. Tang, Q. Fan Sun (Univ. of Waterloo), P. A. Forsyth Sun (Univ. of Waterloo) and J. R. F. McMacken (Univ. of Waterloo), "Performance Issues For Iterative Solvers In Semiconductor Device Simulation," SIAM Journal on Scientific Computing, Vol. 17, No. 1, pp. 100-117, 1996

W.P. Tang, G. Golub (Stanford Univ.), L. C. Huang (Chinese Academy of Science) and H. Simon (Silicon Graphics), "A Fast Solver For Incompressible Navier-Stokes Equations With Finite Difference Methods," SIAM Journal on Scientific Computing, 24 pages, Accepted, 1996

A. Wissink (Univ. of Minnesota), A. Lyrantzis (Purdue), R. Strawn Strawn (US Army AFDD), L. Oliker, and R. Biswas, "Efficient Helicopter Aerodynamic And Aeroacoustic Predications On Parallel Computers," AIAA-96-0153 January 1996.

PAPERS SUBMITTED TO REFEREED JOURNALS

W.P. Tang, T. Chan (UCLA), B. Smith (Argonne National Laboratory) and W. L. Wan (UCLA), "Fast Wavelet Based Sparse Approximate Inverse Preconditioner," BIT. 14 pages, 1996

W.P. Tang, "Effective Sparse Approximate Inverse Preconditioners," SIAM Journal on Matrix Analysis and Applications, 20 pages, 1996

V. SEMINARS AND COLLOQUIA

JAMES J. REUTHER

Workshop for the High Speed Research program to develop a design strategy for a newly released study configuration. NASA Langley Research Center, Langley Virginia, February, 1996.

HSR configuration aerodynamics 1st annual review. NASA Langley Research Center, Langley Virginia, March, 1996.

HSR Technology Configuration Aircraft design review workshop. Seattle Washington, July, 1996.

SIAM Mini-symposium on Optimization Governed by Flow Equations at the 5th SIAM conference on Optimization, Victoria, British Columbia, May 1996.

Computational Aerosciences Workshop, NASA Ames Research Center, Moffett Field, California, August 1996.

WEI- PAI TANG

Workshop on Parallel Unstructured Grid Computations, Argonne National Laboratory, September, 1996.

VI. OTHER ACTIVITIES

MARJORY J. JOHNSON

Co-coordinator for Bay Area Gigabit Network testbed.

Member, U.S. Subcommittee Advisory Group for ISO Technical Committee 20, Subcommittee 13 (ISO/TC 20/SC 13) on Space Data and Information Transfer Systems standards.

Member, NSF review panel for Academic Research Infrastructure Program, June 1996.

Member of several program committees for communications conferences.

JAMES J. REUTHER

Completion of Ph.D. degree from the University of California Davis, Davis, California, June 1996.

Session chairman at the 6th AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Bellevue, Washington, September, 1996.

HSR Certificate of Appreciation, October, 1996.

WEL- PAL TANG

Workshop Organizer, Copper Mountain Iterative Method Conference, Copper Mountain, CO, April, 1996.

VII. RIACS STAFF

ADMINISTRATIVE STAFF

Joseph Oliger, Director - Ph.D., Computer Science, University of Uppsala, Sweden, 1973. Numerical Methods for Partial Differential Equations (03/25/91 - present).

Frances B. Abel, Office and Financial Manager (5/5/88 - 3/8/96).

Consuelo Garza, Administrative Assistant (4/16/96 - present)

Deanna M. Gearhart, Office Manager (2/1/96 - present).
Administrative Assistant II (5/9/88 - 1/31/96).

Rufus White Jr., Systems Administrator (5/17/93 - 2/2/96).

SCIENTIFIC STAFF

Rupak Biswas, Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1991, Large scale scientific computation using parallel and adaptive methods (9/16/91 - 7/31/96).

Dave Gehrt, JD Law, University of Washington, 1972, UNIX system administration, security, and network based tools (1/84 - 7/85, 2/1/88 - present).

Marjory J. Johnson, Ph.D., Mathematics, University of Iowa, 1970, High-performance net working for both space and ground applications (1/9/84 -present).

VISITING SCIENTISTS

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (6/24/96 - 8/23/96).

Tony F. Chan, Ph.D. - Professor of Mathematics, University of California, Los Angeles, Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (8/19/96 - 8/23/96).

Antony Jameson, Ph.D. - McDonnell Professor of Aerospace Engineering, Princeton University, Numerical Methods, computational fluid dynamics, computational sciences (7/1/96 - 9/15/96).

Andrew Sohn, Ph.D. - Assistant Professor, New Jersey Institute of Technology, Dynamic load balancing for grid partitioning on SP-2 (7/17/95 - 8/31/96).

Wei-Pai Tang, Ph.D. - Professor, University of Waterloo, Canada, Numerical solution of partial differential equations, numerical linear algebra, parallel computations (8/1/96 - 8/31/96).

David Zingg, Ph.D. - Associate Professor, University of Toronto, Canada, Development and analysis of high-accuracy numerical methods applicable to simulations of fluid flows, acoustic waves and electromagnetic waves (6/1/96 - 8/21/96)

POST-DOCTORAL SCIENTISTS

James Reuther, Ph.D. - University of California, Davis, numerical optimization aerodynamic shape optimization numerical analysis CFD(4/30/96 - present).

RESEARCH ASSOCIATES

Leonid Oliker, MS - Computer Science, University of Colorado, compilation of data parallel programs (9/1/94 - present).

James Reuther, MS - University of California, Davis, numerical optimization aerodynamic shape optimization numerical analysis CFD(9/6/94 - 4/29/96).

Steven Suhr, MS - Computer Science, Stanford University, programming languages (7/1/92 - present).

RESEARCH ASSISTANTS

Jason Deng - Computer Science, Stanford University, Research for Bay Area Gigabit ATM network testbed, (7/8/96 - 9/13/96).

Jeffrey Townsend, MS- Computer Science, Stanford University, Research for Bay Area Gigabit ATM network testbed, 6/21/95 - 7/31/96).

CONSULTANTS

Marsha Berger, Ph.D. - New York University, Computational fluid dynamics; parallel computing (1/1/93 - present).

Tony F. Chan - Professor of Mathematics, University of California, Los Angeles, Efficient algorithms in large-scale scientific computing, parallel algorithms and computational fluid dynamics (10/01/86 - present).

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956, Global environmental problems and issues (11/1/92 - present).

Pelle Olsson, Ph.D. - Assistant Professor, Stanford University, Initial-boundary value problems for hyperbolic and parabolic PDEs, numerical methods for PDEs on parallel computers (1/4/95 - 8/31/96).

Robert Schnabel, Ph.D. - Professor, University of Colorado, Boulder, Numerical computation especially optimization nonlinear equations, parallel computation (1/1/94 - present).

Wei-Pai Tang, Ph.D. - Professor, University of Waterloo, Canada, Numerical solution of partial differential equations, numerical linear algebra, parallel computations (7/1/94 - present).

Jeffrey Townsend, M.S. - Chromatic, Sunnyvale CA, Research for Bay Area Gigabit ATM network testbed, (8/1/96 - 9/30/96).

